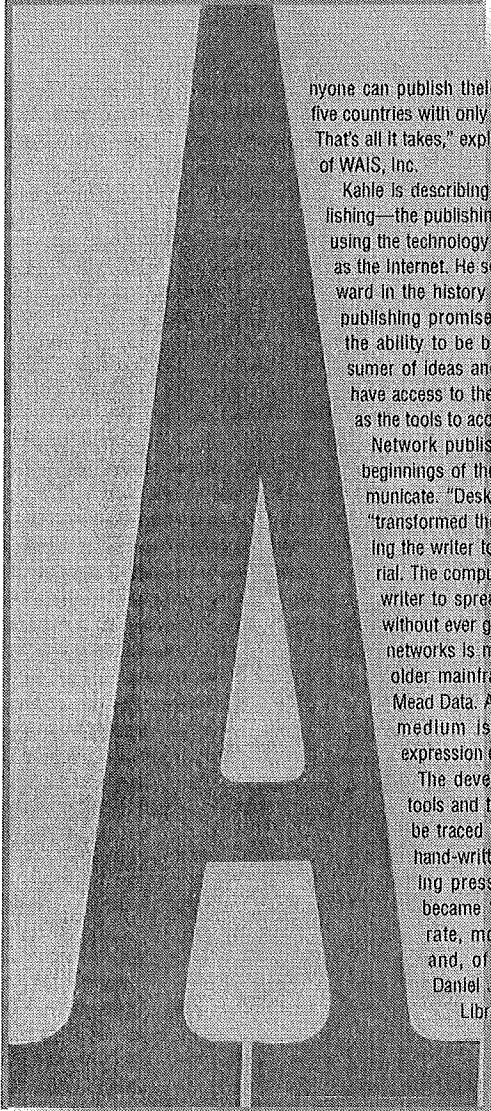


WAIS

—a new vision for publishing

By Michael Robin
Photos By Anne Knudsen



Anyone can publish their "words to people in forty-five countries with only a computer and a telephone. That's all it takes," explains Brewster Kahle, founder of WAIS, Inc.

Kahle is describing what he calls network publishing—the publishing, or dissemination of ideas, using the technology of wide area networks, such as the Internet. He sees it as a quantum leap forward in the history of communication. Network publishing promises to give virtually everyone the ability to be both a publisher and a consumer of ideas and information. Everyone will have access to the tools for publishing as well as the tools to access published information.

Network publishing can be viewed as the beginnings of the next step in how we communicate. "Desktop publishing," says Kahle, "transformed the publishing world by allowing the writer to make 'camera ready' material. The computer networks now allow that writer to spread their words far and wide without ever going to paper... Using these networks is much cheaper than using the older mainframe systems like Dialog or Mead Data. All in all, the new inexpensive medium is creating a new type of expression on the net."

The development of communication tools and their impact on societies can be traced from oral communication to hand-written manuscripts to the printing press. At each step knowledge became "more portable, more accurate, more convenient to refer to, and, of course, more public," as Daniel J. Boorstin, librarian for the Library of Congress, points out in *The Discoverers*. The consequences of making knowledge "more public" is

fundamental to a democratic society. Thomas Carlyle, English essayist and historian, observed in 1863 that:

"He who first shortened the labor of copyists by the device of movable types was disbanding hired armies and cashiering most kings and senates, and creating a whole new democratic world."

The information contained in the early manuscript books was unorganized by today's standards—there were no indexes, punctuation, or even page numbers. The advent of the printing press precipitated an explosion in the amount of knowledge available throughout society, but for many decades there was no useful system for archiving or cataloging these books. All of this combined to make retrieving information difficult.

Similarly, the information available on the Internet today is usually plain, unformatted text, scattered across servers around the world. Searching, retrieving, and sometimes even reading these digital documents often requires familiarity with UNIX commands and other technical jargon. The volumes of information available on desktop computers across the Internet has given renewed meaning to the term "information overload." And until recently, the only way to locate a specific document was to already know exactly where it was—which server, directory, subdirectory, and the file name.

Brewster Kahle is one of the people developing ways to improve access to and retrieval of published documents. First at Thinking Machines and then as founder of San Mateo-based WAIS, Inc., he helped create WAIS, Wide Area Information Servers, a remote search and retrieval resource for networks such as the Internet. WAIS is one of several relatively new resources on the Internet that seeks to help users navigate online to provide and/or find information. Kahle hopes that WAIS will help lay the foundation for transforming the way we communicate.

I met with Kahle to discuss what WAIS software is, how it works, WAIS, Inc., and his vision of network publishing.



As more and more people are getting on networks and large corporations are building private enterprise networks, better tools for navigating and searching the nets are becoming available, tools such as gopher, WAIS,archie, Web, etc. Putting WAIS in context, how does it fit into this puzzle?

WAIS is one of the tools that are being used now on the Internet. There are three really wildly popular ones—gopher, World Wide Web, and WAIS. If you think about this as what we're constructing is a large network book, gopher is the table of contents, a hierarchical browsing approach. World Wide Web is hypertext pages. It's the realization of Ted Nelson's dream, in terms of technology, of clicking from one page to another, jumping to author to author to author. WAIS is the back of the book, the index. It's the place you turn when you know what you want. WAIS is trying to help people search through gigabytes of information that's located in hundreds of locations around the world.

To use gopher you don't need to know particular sites or where the information is on the nets, but for WAIS you need to know specific locations.

It's a two-step process with WAIS. One, you have to find the data source that you want to look through. There are directories that help you find the right data collections, whether it's a poetry server or weather server, whether it's a genetics database in Singapore or the IBM PC frequently asked questions server in Mexico City; you have to find which of those you want to use.

You ask the directory of servers a question. Say you're interested in how to hook up your IBM PC to an Ethernet network. You'd ask the question, "I'm trying to hook up my IBM PC to an Ethernet network," to the directory of servers. Hopefully there will be some descriptions of databases that have some of those words and phrases in them and they'll be suggested back to you. There may be four or five databases that will have information on IBM PCs.

You would then select one or several of them and pose your question to those databases, which can be located anywhere around the world. You don't care where the servers are. You use a server the same way you'd use it as if the data were on your own local machine, and it all works at pretty much the same speed. You can search through gigabytes in just a few seconds no matter where it is in the world.

The idea is that you actually ask your question in the English language, like the example, "I'm trying to hook up my IBM PC to an Ethernet network." The servers don't understand what you're saying. They're just trying to find documents that have those words and phrases in them. If the words and phrases exist exactly, as in your question, the document receives more weight. If it's in the headline, it receives even more weight. The list of documents are ranked and the top ones come back at the top of the list. The server says you might want to look at these. You can browse them, then by clicking on them, retrieve the information you want.

You can also say, "I like that one. Find me more like that one." WAIS uses the server again by getting this positive feedback from the user. It knows a lot more what it is you are interested in. It can use a document as a big search term to find other similar documents.

What is the Directory of Servers?

The WAIS software indexes all the words and phrases in each document so that

you can find appropriate documents. A "Directory of Servers" is a database of descriptions of servers and a major one is operated by Thinking Machines. This can help find appropriate servers to use.

We hope that the combination will help people find the most appropriate documents to read.

From the end-user perspective, what is happening when I do a WAIS search?

WAIS is the system that sits next to the publisher making their information available. The users can use different interfaces to get to the information—gopher users work with certain interfaces, World Wide Web users will use something else. The America Online interface to WAIS is going to be available very soon. So there are lots of different interfaces to the same information.

If you're cruising the net using Web and hit a search box, often you're using a WAIS resource behind it. If you're using gopher and you hit a question mark icon, and up pops a little box, "What do you want to search?" you're probably using WAIS underneath.

The idea is that publishers have something to say to lots of different audiences and those audiences are going to use whatever devices they want for finding that information. Everybody's got a religious bent, whether they want their PC, their Macintosh, their Newton, their General Magic machine—whatever interface they want to use, we just want to be able to have the information get to those users.

Have there been conflict or problems by the ways people in different professions or cultures ask questions?

Some people are very careful about how they ask things and some people just blather. There are some people that say "IBM PC" or "Ethernet transceiver" and that's all that the server gets to try to figure out from the tens of thousands of documents that it may have which ones you want. Some people will use English language, as if they're talking to another person. We're trying to make a system that will work in those environments.

Another class of users are the real expert searchers, the ones that know how to use Boolean language and fielded search, they know this word, or I want this word next to this word and the author equals this; these are the librarians and professional searchers. Lawyers also know these types of things. We're trying to make a system that is useful by everyday people to just browse and have fun, and for other people that really know what they want to be able to screw down and get just those three documents.

Is WAIS available for different languages?

Yes, the European languages, and Fujitsu is making a Japanese version.

I saw an interesting thing that Fujitsu had done. I don't know if they're going to release it, but they integrated online machine translation so you can ask a question in English against a Japanese database and it translates the Japanese documents into English for you to see. It was mind-blowing. Of course it's not perfect translation. Even people aren't very good at translating English to Japanese or vice versa, but it did make it readable. The aspect of making it so that people in Japan can communicate more freely with people in English through these publishing mediums with electronic assistance is astounding.

There is an important step before full-text online translations, and that

is just being able to search across servers in the different languages. For instance, posing an English question to a Spanish-language database.

You can do that now and the standards are catching up to make that done in a completely standard way. We still have problems of how do you represent all of these languages in computer-readable form? Kanji is a problem, not just that it doesn't fit into 8-bit ASCII. It takes 16 bits or more. But trying to find where words start and stop in Kanji is a challenge. There are many aspects to WAIS that aren't just of displaying document problems. WAIS really has to understand a little bit about the document to know where the words and phrases are and what's important about them so you don't get irrelevant material. That changes from culture to culture and language to language.

In addition to text, how does WAIS deal with multimedia documents?

You can search databases of text as well as images and video clips or whatever. WAIS uses tags—natural language tags which are descriptions of documents, descriptions of pictures, or maybe it's the sound track for the video. Define what video you want and then, when you double-click on it, it can be anything, a word processor document, pictures of Europe, rock music recordings.

And retrieving them?

It can be done over the networks transparently. There's a speed issue. Text is able to be done easily over modems. For still images, modems are getting to be a little slow, but the Internet connections, which tend to be 56K, are fine for this. We need faster speeds for doing video and audio.

There's no reason to think that this technology is not going to help make video archives or video artists able to publish their media where they are unable to now through the Blockbuster-type chains. It's changing the publishing equation by allowing individuals to publish on the networks as opposed to having to go through large-scale publishing enterprises.

WAIS, Inc. provides the back-end, or software for the publisher, while the front-end software is freeware?

Most of the WAIS software is not done by WAIS, Inc. It's been done by the freeware community, which is this phenomenally rich, interesting group of people that are spanning the globe working together to make an open information infrastructure.

Most of the user interfaces that are available now are available for free on the Internet. You can use file transfer from wais.com to get them. There's even free server software, but for those people that want to go into production, that want to have really good searching, they tend to want a commercial-grade tool. That's what we sell. That's how we support ourselves. But we're very much dependent on the freeware world to keep the freeware good and also, the most important thing is to get rich information resources out there in lots of specialties.

And how closely do you work with the freeware people?

Weekly basis, daily basis. We're a major distribution site for the freeware. We put out some of our enhancements as freeware. Basically, we need critical mass in terms of a publishing system so we're all talking the same protocol. The key piece is Z39.50, URLs, and a bunch of other acronyms out there. The idea is that it's open protocol, not a protocol that is "open" that's really owned by a company.

Do you foresee some of the front-end software being commercialized?


Absolutely. Apple has just announced that they are "gatewaying" from their AppleSearch product to the Internet WAIS servers. So you can, from your Macintosh environment, search the Internet using this product that's due out later this year.

What do you mean by "gatewaying" from AppleSearch to WAIS? Does that mean it is another type of WAIS front-end?

Apple doesn't like to think of it that way. I would put it the other way around—people in the AppleSearch environment can get to WAIS resources. So people in the America Online community can get to WAIS resources, like people in the gopher community can get to WAIS resources.

What is the relationship of AppleSearch to the WAIS engine?

AppleSearch will be able to search WAIS resources on the Internet. Apple is one of the original members of the WAIS project, which included Apple, Dow Jones, Peat Marwick, and Thinking Machines. An existing project at Apple was working on a similar problem for the local area networks that became AppleSearch. WAIS offered a method to opening it to the outside world.



WAIS In The Market

We talked about the end-user. What is the provider or publisher doing with WAIS?

There's the traditional publishers who are using this as a mechanism to distribute their work in a new way. For examples of some of those involved, Dow Jones will be distributing the *Wall Street Journal* and same day *New York Times* on the Internet through the WAIS system. Another is Encyclopedia Britannica, who is bringing their whole encyclopedia to the Internet through WAIS and World Wide Web. This is where you get traditional publishers seeing WAIS as a mechanism to get to people all over the world by having their information in one place and selling subscription services.

We're also seeing lots of other people becoming publishers. And that's the new thing I see that is societally interesting, as new things happen because we have this new technology. So, for instance, Sun Microsystems is distributing all of its information, press releases, bug fixes, etc. over the nets using WAIS. Small groups, like the A Right To Keep And Bear Arms group, have also set up their servers. It's not my particular cause but I'm really glad that they're able to have a mechanism to make their point of view known. So we're getting lots of people publishing lots of different types of documents. We want musicians to be able to not only publish for free but also be compensated for their work where before they were not able to find the niche within the traditional publishing environment.

WAIS and WAIS, Inc. grew out of something you were working on at Thinking Machines. Unlike most other Internet tools, there seems to have been a commercial vision for this project from the get-go. What were you doing at the time that led to this project?

Thinking Machines is a parallel computing company. In the early '80s it sud-

Who are the primary markets for WAIS?

We have two major ones and two minor ones. One is the government. The Environmental Protection Agency is making a set of services available on the net through WAIS. The Library of Congress is making some of its picture collections and things they had on CD-ROM available on the net. So the government has a traditional role of distributing information to very wide populations.

Another is publishers, traditional publishers that want to use this in a new way. We have two others, which are large corporations that want to publish internally and to the outside world. As we have more globally distributed companies, keeping people in touch about what is going on within the company is a real challenge.

Libraries are also trying to figure out their role in the new worlds of networks and electronic text.

Are there commercial WAIS servers?

They're just starting. There's a commercial WAIS server that's offering some government information. You'll see the roll-out of the *Wall Street Journal* in April. *Encyclopedia Britannica* is in test now and it will go live in the fall. So it's all just happening now. It's very exciting to see the shifts happening, there are a number of commercial enterprises setting up shop where they're selling information over the net.

Are there mechanisms for payment?

Network publishing can support any type. You can charge per document, per search, however you desire. The mechanism people are looking at most is subscription-based, you can use the *Wall Street Journal* and all the back issues, all you want for a month for similar pricing to what you'd pay for a paper copy.

That's the thought that's going into pricing in these environments. It's low-price so end-users, everyday people, can use these information sources and they can use them all they want.

Besides the projects by Dow Jones and Encyclopedia Britannica, are large corporations using WAIS on their enterprise networks?

Perot Systems and Lockheed are examples. Perot Systems put all the resumes of everybody in the company on its computer so you can find people. They have contract proposals, presentations, and have also downloaded some CD-ROMs, by paying the CD-ROM publisher and making it available to their company.

Is there competition to WAIS? Others like Mead Data, Lexis/Nexis? Is one of the major differences between those services and WAIS the issue of open protocols?

The difference between us and the Dialogs and Mead Datas, which are centralized publishing models, is that WAIS is decentralized. It's uncontrolled and uncontrollable. Anybody can go and use the software and make their words known. So it's more based for open networks and distributed computers. Lotus Notes has been targeted mostly for LAN-based environments within a company where WAIS was designed for cruising databases all over the world. You don't know exactly what you're looking for, so it's oriented a little bit differently, although all of these systems, we hope, will become compatible with using WAIS resources. So we would love Dialog's and Mead Data's data to be available on the Internet through the open protocols of WAIS, as well as Lotus Notes. We want those users to be able to get at WAIS resources. Are either of these happening? There've been talks about it but nothing concrete.

I don't think that there's really direct competition because the networks are too new. Gopher, World Wide Web, WAIS, we all gateway to each other. They're all based on open protocols. I'd say the way we'd lose is if we didn't do a system that's good enough so that there's room for a proprietary solution to come in. That's the danger. All the other companies, the search engine companies, the database companies, got very excited because they can get at more users, they can get at more information, that's a win-win situation.

denly had hundreds of times more computing power than people ever had before. The question was, What do you do with it? Well, you can simulate weather better. You could try to find oil under the ground. But we also thought that there was something we could do that would be usable by everyday people. Finding the right things to read was the one that we had in our minds from a project that we began back at MIT.

When the [Thinking Machines] computer came along, we tried out search-

ing through gigabytes. In 1985, we did a project searching through 15GB with a twenty thousand-term query and it only took three minutes. It took a supercomputer to do it, and this was one hundred times faster than anything else at the time. You could browse through colossal collections. Of course 15GB is not that much anymore; that's about what you work with on a workstation or high-end PC.

What Thinking Machines had was a view of the future, of what things would look like in ten or fifteen years. Now we're at that time and it doesn't require a supercomputer. In the mid-'80s, when Dow Jones bought one of these systems to search through 450 magazines and newspapers to find the articles you wanted required a supercomputer. Now it can be done on a Sun microcomputer or almost any kind of UNIX box.

Another way to look at where the project came from is, "yes, it came from a very commercial background." We think it's important that people be able to be compensated for their work, or you can't have an enduring environment. When the printing press came along there wasn't the concept of copyright, and it took 150 years to get the royalty systems together that we now have for books and newspaper publishing. At that time writers more or less donated their work, or got a fixed fee, a one-time fee for their work. If we can help the network environment establish a process so people can be compensated for their work, the whole field will expand much more quickly. We want high-quality information as well as free information out there. The technology is not innately worth anything. It's the content that you can get to people that's important.

What brought Thinking Machines, Dow Jones, Apple, and Peat Marwick together to develop WAIS?

Thinking Machines and Dow Jones built an innovative search system called DowQuest, and we thought it would radically change the world. Well, the world wasn't terribly different. Yes, they made money on it. It was an interesting product. But it didn't affect everyday people. The question was, Why? Was the pricing wrong? Were the networks not there? Did it need a graphic user interface? What were the pieces?

So I spearheaded a project to say "all right, let's figure it out. Let's go and get a group of companies to work together in relative secrecy to figure out what it takes" Apple Computer is world-renowned for great user interface design; Thinking Machines for search engines; Dow Jones for one-stop shopping for information for business people; and Peat Marwick represented a community that knew what their time was worth—they were a perfect community, non-techies. If we could make it usable by the partners at Peat Marwick, we'd have a system that would work. And we built a system in nine months, a crack team of people across the country. We found that, yes, Peat Marwick's people did like it, they would use day by day.

The problems were the networks. The network was just too hard to construct. This was back in 1989. The Internet was still a research network. It really wasn't mainstream at all. And that's when we looked around and said, "Well, the Internet is working. Let's use the Internet as our model."

Thinking Machines produced a freeware release. It was a little hard to argue for this with the management at Thinking Machines. We were saying, "We want to take this new idea and distribute it in the public domain, no copyrights, no patents, no control. Anybody can take it, copy it, merge it into their own programs."

"Why would you possibly want to do that?" was a question I had to answer. And the answer was that we wanted to catalyze a market for information servers. We needed a critical mass and the only way to do it was to seed the

To find out more information about WAIS,

You can send electronic email to info@wais.com. There is free software available by anonymous ftp from wais.com and you can telnet for a dumb terminal interface to WAIS, telnet to wais.com; and login as "wais". You can also use WAIS through World Wide Web/Mosaic by <http://wais.com/>. WAIS, Inc. in San Mateo can be reached at (415)327-WAIS.

market with a good enough system to get the ball rolling. Thinking Machines is a long-range company and they said, "Great, let's do it." They wanted, after the market was built up, to sell supercomputers to this market. They've sold a couple into it already, but as the market grows, there will be colossal text collections that will need their computers. So it was not money badly spent from Thinking Machines' point of view. But it did set a precedent that these systems are going to be open.

Does WAIS offer a new business model for the networks? What is the relationship of WAIS, Inc. to the Internet culture where, traditionally, information was distributed freely? Also, what would stop anyone else from taking a freeware version, enhancing it, and providing the support services you offer?

That's several questions. Is this a new business model for how to disseminate information? Absolutely.

I think there will be an enduring need for free components—they will always be available, but may not have all the features. There are people using this system in schools that would just not be able to buy much software. So, it's important to have a free version out there.

It's also important that when people need more features or capability that there's an avenue for them to get what they need, which is often not the case on the Internet. You can't buy quality versions of some of the things on the Internet even if you wanted to, yet. We think that there is room, based on open systems, to have commercial and free versions at the same time.

You asked the question, "Can someone come in and compete with us?" Absolutely. We're playing the open systems game. When we started the project with Apple, Dow Jones, Thinking Machines, we said it was going to be based on open protocols. In fact, at that time the protocols weren't very good; we needed to get them better. We've spent about half of our engineering resources

making the open protocols good enough for our competitors to use. Why? Because it needs to be an open system to expand, to be an interesting environment.

As a business model, providing commercial support for freeware, how similar is WAIS, Inc. to Cygnus, which supports GNU software?

We're quite different from the business model of Cygnus. Our server has been rewritten from scratch. I helped write some of the freeware version of the server software. But when we formed the company, we started over and rewrote the system from scratch based on the protocol Z39.50. It's a completely different implementation that's much higher quality, and it's portable and has been extended in many different environments. Unlike Cygnus, which is all the same code base and they provide support, this is a different product. There are other companies implementing WAIS standards and those are completely different implementations. The important thing is that we all have the same protocol. So where most people talk the same code base in the PC world, in the network environment it's not the code that's important, it's the protocols.

So that everything, whether it's a WAIS server from you or something developed from freeware, they all communicate.


They all communicate.

You have written about network publishing, describing it not as something replacing books or even saving trees, but as a new media form. What is your vision of network publishing?

For me the important aspect of network publishing is not that you can now peruse the library in Moscow from your desktop machine. It's turning the equation around.

It's that anybody with something to say can now have a forum to say it and they can make it available in this new and open way.

Where successful technologies have



continued from page 86

opened up more and more ways to make their words known, network publishing is a huge jump. You can now publish your words to people in forty-five countries with only a computer and a telephone. That's all it takes.

There are very few generations that get to see the development of a new technology for how people communicate. That generation gets to see all sorts of wild things happen: industries come and go, what it means to be in a family and how companies work changes. All of those things change dramatically based on a new communication technology. I'd suggest network publishing is such a change. It means that you don't have to go through the established hierarchies that have built themselves up around older technologies of information distribution. People can take their photographs of Asia and make it available on the networks and find other people that have similar interests. You can take your MIDI files of music recordings and make those available, and find other people that are interested in your kind of music; and, as we're seeing now, people are starting to do it for pay, not just for free. It's more than just electronic mail or bulletin boards. It's not just conversational. People are putting together real works that are composed and created to work over time. They're not just a snapshot of email.

Is electronic network publishing changing our notion of what a book or document is? It opens the possibilities for customized books, or books that change over time, or books that can be easily read in a nonlinear fashion.

Yes. We're only starting to see some of the new types of literature that will come out of this new medium. As [Marshall] McLuhan wrote, the new medium contains the old medium. So the first thing you do in a new medium is just make a copy of the old medium. Later it will start to evolve on its own. It's hard to know where it's all going, but one way to try to think about it is to look at the different types of books we have. Try to think what would make a really great electronic encyclopedia? Well, it wouldn't just link to the documents within itself, it would also point out to the real world and point to current newspaper articles. Or, look at an atlas. You wouldn't want a static picture of each map. You'd want to be able to zoom in and out, move around, see it over history, and turn different knobs to be able to interact with maps in a new and different way. We'll see every book type and every information type evolve as the market grows up on the nets. The key piece is to make sure there's a market and not just the technology.

One vision for the future is that publishing houses will become a company with a server, providing an outlet for the writer or photographer or musician who does not have their own server.

Yes. Right now it takes quite a bit of technical sophistication to make yourself an Internet node and run these services, but it's becoming easier and easier. There are already service bureaus which make somebody else's catalog available on the Internet through WAIS or World Wide Web.

It will get to the point where small companies or even an individual's machine will be able to be those nodes. It's becoming easier and easier to run these small-scale printing presses for Internet distribution.

What are some of the other issues network publishing will force us to reexamine?

The world is a dynamically changing environment. The archivists have got some real work to do. I am working with a group of people founding the Library of Alexandria Foundation, which is trying to archive some of the works that are created on the net that were never meant to be printed.

Every librarian has two hats. There's the access hat and the archiving hat. Right now on the Internet we're very heavily biased on the access, but there's cultural changes and sociological changes going on that are not even being documented, they're not being saved, and it's time to start addressing these issues.

Is the question of authenticity of a document part of archiving—being able to make sure the copy you have is made from an unaltered original?

Authenticity on the networks, security, and identity on the networks are all very real problems, and haven't been resolved very well. There are people going onto and making money on these systems that are actually very easy to spoof. So, it's possible to steal on the networks as real commercial enterprises join in. And

those access methods really need to be moved towards actually figuring out who is a person and is their money good? Is their Visa card number good or is it just stolen? Those sorts of aspects are still to come.

And digital signatures on documents?

Digital signatures is another technology for authenticating a document. It's becoming less of a problem. If you want to get a copy you go back to the original source, so you don't necessarily have to have many copies of a particular document floating around on the net. You can just reference the original copy on someone's hard drive and when you want it, click on it, and you retrieve the original document one more time from the original location.

In that way it's different from the printing environment where there are copies floating around. They can be altered. In the network publishing model, the publishers control the distribution of their work and they certify the authenticity of the work that they have. That's what a publisher does.

Archiving for the nets reminds me of the evolution of printed material and how systems needed to be developed to manage the growing volumes of information, documents, and books that were being collected. How will we handle all the new information

available electronically?

Let me try to answer this by saying that when people are complaining of information overload it means the tools aren't good enough to find only those things that they want to read. And the tools aren't good enough, yet.

We think that making the tools better for searching will include using "editors" to help select works and to let you know that these are the hot articles. You might subscribe to several different editors to help select the different articles that you want to read. Editors will be

the next wave of navigation tools on wide-area networks.

The other question is archiving. This is just starting and there is nothing good to say, yet. In terms of actual experience, there are only horror stories of archiving in the electronic publishing world. But what will help on the archiving side is that it is really cheap to archive things. With an \$8 tape you can store 5GB.

When the printing press came along there wasn't the concept of copyright and it took 150 years to get the royalty systems together that we now have for books and newspaper publishing.

Is copyright another issue with archiving?

I would say copyright is more an issue of access, restricting access, as opposed to the archiving rule which is how to save it for the future.

We are beginning to see the early battles over electronic copyright, what rights are being bought and sold, and how authors or creators should be compensated.

Yes. Exactly how the compensation structure will work in the electronic world has not been figured out and it will take time. The key piece is that somebody will start making some money somewhere. That's what we are trying to figure out now and we're in the very beginning throes.

We're excited that successful models of putting across content that people will want to pay for on the net, open networks, is starting to happen. All the rest can figure itself out afterwards. It's an important question, but if you don't answer the first one, we don't even have a game to play. ■